

FORM PTO-1390
(REV. 9-2001)

U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE

ATTORNEY'S DOCKET NUMBER

TRANSMITTAL LETTER TO THE UNITED STATES
DESIGNATED/ELECTED OFFICE (DO/EO/US)
CONCERNING A FILING UNDER 35 U.S.C. 371

10002166-2

U.S. APPLICATION NO. (If known, see 37 CFR 1.5

09/980920

INTERNATIONAL APPLICATION NO.

PCT/US00/14282

INTERNATIONAL FILING DATE

24 May 2000

PRIORITY DATE CLAIMED

24 May 1999

TITLE OF INVENTION

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

APPLICANT(S) FOR DO/EO/US

KRAUSE, Michael; RECIO, Renato, John

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

1. ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. ☐ This is an express request to begin national examination procedures (35 U.S.C. 371(f)). The submission must include items (5), (6), (9) and (21) indicated below.
4. ☐ The US has been elected by the expiration of 19 months from the priority date (Article 31).
5. ☒ A copy of the International Application as filed (35 U.S.C. 371(c)(2))
 - a. ☐ is attached hereto (required only if not communicated by the International Bureau).
 - b. ☐ has been communicated by the International Bureau.
 - c. ☒ is not required, as the application was filed in the United States Receiving Office (RO/US), but is attached hereto.
6. ☐ An English language translation of the International Application as filed (35 U.S.C. 371(c)(2)).
 - a. ☐ is attached hereto.
 - b. ☐ has been previously submitted under 35 U.S.C. 154(d)(4).
7. ☐ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3))
 - a. ☐ are attached hereto (required only if not communicated by the International Bureau).
 - b. ☐ have been communicated by the International Bureau.
 - c. ☐ have not been made; however, the time limit for making such amendments has NOT expired.
 - d. ☐ have not been made and will not be made.
8. ☐ An English language translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371 (c)(3)).
9. ☐ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. ☐ An English language translation of the annexes of the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11 to 20 below concern document(s) or information included:

11. ☐ An Information Disclosure Statement under 37 CFR 1.97 and 1.98.
12. ☐ An assignment document for recording. A separate cover sheet in compliance with 37 CFR 3.28 and 3.31 is included.
13. ☒ A FIRST preliminary amendment.
14. ☐ A SECOND or SUBSEQUENT preliminary amendment.
15. ☐ A substitute specification.
16. ☐ A change of power of attorney and/or address letter.
17. ☐ A computer-readable form of the sequence listing in accordance with PCT Rule 13ter.2 and 35 U.S.C. 1.821 - 1.825.
18. ☐ A second copy of the published international application under 35 U.S.C. 154(d)(4).
19. ☐ A second copy of the English language translation of the international application under 35 U.S.C. 154(d)(4).
20. ☐ Other items or information:

FORM PTO-1390 (REV 9-2001) page 2 of 2

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

Commissioner for Patents
Washington, D.C. 20231

PRELIMINARY AMENDMENT

Sir/Madam:

This Preliminary Amendment modifies the Utility Patent Application filed herewith.

Please amend the above-identified application as follows:

IN THE CLAIMS

The following is a clean version of the entire set of pending claims, where claim 1 remains as previously filed and new claims 2-34 are added with this Preliminary Amendment.

1. A method of managing memory in a distributed computer system, the method comprising:
 - binding a remote key to a first address representing a contiguous memory address range at a first endnode;
 - sending the bound remote key and first address from the first endnode to a second endnode on a communication fabric; and
 - performing a remote direct memory access operation from the second endnode to access the contiguous memory address range including sending the bound remote key and the first address from the second endnode to the first endnode on the communication fabric.
2. The method of claim 1 further comprising:
 - controlling local memory access protection in the first endnode with a virtual memory manager in an operating system kernel process.
3. The method of claim 1 further comprising:
 - comparing the bound remote key and the corresponding first address supplied by the second endnode to the bound remote key and corresponding first address in the first endnode.

2017-10-02 09:55:50

Preliminary Amendment

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

4. The method of claim 3 wherein if the bound remote key and corresponding first address supplied by the second endnode do not match the bound remote key and first address in the first endnode, the second endnode is not granted access to the contiguous memory address range.
5. The method of claim 3 wherein if the contiguous memory address range represented by the first address bound to the remote key supplied by the second endnode is invalid, the second endnode is not granted access to the contiguous memory address range.
6. The method of claim 1 wherein the first address is an effective address pointing to an address space in memory accessible by a consumer process.
7. The method of claim 6 wherein the effective address points to a virtual address space.
8. The method of claim 7 further comprising:
 - comparing the bound remote key and the corresponding first virtual address supplied by the second endnode to the bound remote key and corresponding first virtual address in the first endnode; and
 - handling a page fault condition in the first endnode caused by the first virtual address bound to the remote key supplied by the second endnode not being previously mapped by an operating system of the first endnode.
9. The method of claim 6 wherein the consumer process is a user process.
10. The method of claim 6 wherein the consumer process is a kernel process.
11. The method of claim 1 wherein the first address is a virtual address accessible by a consumer kernel process.
12. The method of claim 1 wherein the binding includes associating the first address to the remote key with a consumer process employing a bind remote key verb.

Preliminary Amendment

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

13. The method of claim 1 further comprising:
obtaining at least one remote key with a consumer process employing an allocate remote key verb.
14. The method of claim 1 further comprising:
unbinding the remote key from the first address with a consumer process employing an unbind remote key verb.
15. The method of claim 13 further comprising:
retiring at least one remote key that was previously obtained via the allocate remote key verb with the consumer process employing a deallocate remote key verb.
16. The method of claim 1 wherein the remote key cannot be used to protect more than one memory region at a given instant.
17. The method of claim 1 further comprising:
reusing the remote key after the remote direct memory access operation from the second endnode is completed.
18. The method of claim 1 further comprising:
disabling a translation for the remote key after the remote key is used for the remote direct memory access operation from the second endnode.
19. A distributed computer system comprising:
communication fabric;
a first endnode including:
a first consumer process which binds a remote key to a first address
representing a contiguous memory address range accessible by the consumer process;
and
a first interface controller sending the bound remote key and first address from the first endnode on the communication fabric; and

Preliminary Amendment

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

a second endnode including:

a second interface controller receiving the bound remote key and first address from the first endnode via the communication fabric; and

a second consumer process performing a remote direct memory access operation from the second endnode to access the contiguous memory address range including sending, via the second interface controller, the bound remote key and the first address from the second endnode to the first endnode on the communication fabric.

20. The distributed computer system of claim 19 wherein first endnode further includes: an operating system kernel process having a virtual memory manager controlling local memory access protection in the first endnode.
21. The distributed computer system of claim 19 wherein in the remote direct memory operation if the bound remote key and corresponding first address supplied by the second endnode do not match the bound remote key and first address in the first endnode, the second endnode is not granted access to the contiguous memory address range.
22. The distributed computer system of claim 19 wherein in the remote direct memory operation if the contiguous memory address range represented by the first address bound to the remote key supplied by the second endnode is invalid, the second endnode is not granted access to the contiguous memory address range.
23. The distributed computer system of claim 19 wherein the first address is an effective address pointing to an address space in memory accessible by the consumer process.
24. The distributed computer system of claim 23 wherein the effective address points to a virtual address space.
25. The distributed computer system of claim 19 wherein the consumer process is a user process.

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

- | | 1980 | 1981 | 1982 | 1983 | 1984 | 1985 | 1986 | 1987 | 1988 | 1989 | 1990 | 1991 | 1992 | 1993 | 1994 | 1995 | 1996 | 1997 | 1998 | 1999 | 2000 | 2001 | 2002 | 2003 | 2004 | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 | 2018 | 2019 | 2020 | 2021 | 2022 | 2023 | 2024 | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | 2031 | 2032 | 2033 | 2034 | 2035 | 2036 | 2037 | 2038 | 2039 | 2040 | 2041 | 2042 | 2043 | 2044 | 2045 | 2046 | 2047 | 2048 | 2049 | 2050 | 2051 | 2052 | 2053 | 2054 | 2055 | 2056 | 2057 | 2058 | 2059 | 2060 | 2061 | 2062 | 2063 | 2064 | 2065 | 2066 | 2067 | 2068 | 2069 | 2070 | 2071 | 2072 | 2073 | 2074 | 2075 | 2076 | 2077 | 2078 | 2079 | 2080 | 2081 | 2082 | 2083 | 2084 | 2085 | 2086 | 2087 | 2088 | 2089 | 2090 | 2091 | 2092 | 2093 | 2094 | 2095 | 2096 | 2097 | 2098 | 2099 | 2100 | 2101 | 2102 | 2103 | 2104 | 2105 | 2106 | 2107 | 2108 | 2109 | 2110 | 2111 | 2112 | 2113 | 2114 | 2115 | 2116 | 2117 | 2118 | 2119 | 2120 | 2121 | 2122 | 2123 | 2124 | 2125 | 2126 | 2127 | 2128 | 2129 | 2130 | 2131 | 2132 | 2133 | 2134 | 2135 | 2136 | 2137 | 2138 | 2139 | 2140 | 2141 | 2142 | 2143 | 2144 | 2145 | 2146 | 2147 | 2148 | 2149 | 2150 | 2151 | 2152 | 2153 | 2154 | 2155 | 2156 | 2157 | 2158 | 2159 | 2160 | 2161 | 2162 | 2163 | 2164 | 2165 | 2166 | 2167 | 2168 | 2169 | 2170 | 2171 | 2172 | 2173 | 2174 | 2175 | 2176 | 2177 | 2178 | 2179 | 2180 | 2181 | 2182 | 2183 | 2184 | 2185 | 2186 | 2187 | 2188 | 2189 | 2190 | 2191 | 2192 | 2193 | 2194 | 2195 | 2196 | 2197 | 2198 | 2199 | 2200 | 2201 | 2202 | 2203 | 2204 | 2205 | 2206 | 2207 | 2208 | 2209 | 2210 | 2211 | 2212 | 2213 | 2214 | 2215 | 2216 | 2217 | 2218 | 2219 | 2220 | 2221 | 2222 | 2223 | 2224 | 2225 | 2226 | 2227 | 2228 | 2229 | 2230 | 2231 | 2232 | 2233 | 2234 | 2235 | 2236 | 2237 | 2238 | 2239 | 2240 | 2241 | 2242 | 2243 | 2244 | 2245 | 2246 | 2247 | 2248 | 2249 | 2250 | 2251 | 2252 | 2253 | 2254 | 2255 | 2256 | 2257 | 2258 | 2259 | 2260 | 2261 | 2262 | 2263 | 2264 | 2265 | 2266 | 2267 | 2268 | 2269 | 2270 | 2271 | 2272 | 2273 | 2274 | 2275 | 2276 | 2277 | 2278 | 2279 | 2280 | 2281 | 2282 | 2283 | 2284 | 2285 | 2286 | 2287 | 2288 | 2289 | 2290 | 2291 | 2292 | 2293 | 2294 | 2295 | 2296 | 2297 | 2298 | 2299 | 2300 | 2301 | 2302 | 2303 | 2304 | 2305 | 2306 | 2307 | 2308 | 2309 | 2310 | 2311 | 2312 | 2313 | 2314 | 2315 | 2316 | 2317 | 2318 | 2319 | 2320 | 2321 | 2322 | 2323 | 2324 | 2325 | 2326 | 2327 | 2328 | 2329 | 2330 | 2331 | 2332 | 2333 | 2334 | 2335 | 2336 | 2337 | 2338 | 2339 | 2340 | 2341 | 2342 | 2343 | 2344 | 2345 | 2346 | 2347 | 2348 | 2349 | 2350 | 2351 | 2352 | 2353 | 2354 | 2355 | 2356 | 2357 | 2358 | 2359 | 2360 | 2361 | 2362 | 2363 | 2364 | 2365 | 2366 | 2367 | 2368 | 2369 | 2370 | 2371 | 2372 | 2373 | 2374 | 2375 | 2376 | 2377 | 2378 | 2379 | 2380 | 2381 | 2382 | 2383 | 2384 | 2385 | 2386 | 2387 | 2388 | 2389 | 2390 | 2391 | 2392 | 2393 | 2394 | 2395 | 2396 | 2397 | 2398 | 2399 | 2400 | 2401 | 2402 | 2403 | 2404 | 2405 | 2406 | 2407 | 2408 | 2409 | 2410 | 2411 | 2412 | 2413 | 2414 | 2415 | 2416 | 2417 | 2418 | 2419 | 2420 | 2421 | 2422 | 2423 | 2424 | 2425 | 2426 | 2427 | 2428 | 2429 | 2430 | 2431 | 2432 | 2 |
|--|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|---|
|--|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|---|

Preliminary Amendment

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

REMARKS

This Preliminary Amendment modifies the above identified Application filed herewith. In this Preliminary Amendment claims 2-34 have been added. **Applicants request that filing fees be calculated on the basis of entry of this Preliminary Amendment.**

Claims 1-34 remain in the application and are presented for consideration and allowance.

It is believed that all claims 1-34 of this application are in condition for allowance. A notice to that effect is respectfully requested.

2011-10-02 09:02:50

Preliminary Amendment

Applicant: Michael R. Krause et al.

Filed: Herewith

Docket No.: 10002166-2

Title: MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

Any inquiry regarding this Preliminary Amendment should be directed to either Patrick G. Billig, Esq. at Telephone No. (612) 573-2003, Facsimile No. (612) 573-2005 or William J. Streeter, Esq. at Telephone No. (970) 898-3886, Facsimile No. (970) 898-7247. In addition, all correspondence should continue to be directed to the following address:

Hewlett-Packard Company
Intellectual Property Administration
P.O. Box 272400
3404 E. Harmony Road, M/S 35
Fort Collins, Colorado 80527-2400

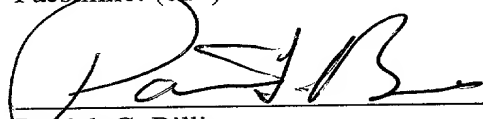
Respectfully submitted,

Michael R. Krause et al.

By their attorneys,

DICKE, BILLIG & CZAJA, P.A.
701 Building, Suite 1250
701 Fourth Avenue South
Minneapolis, MN 55415
Telephone: (612) 573-2003
Facsimile: (612) 573-2005

Date: 11-26-01
PGB:cmj

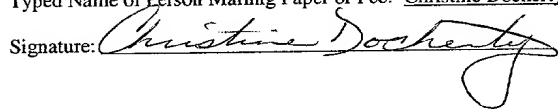

Patrick G. Billig
Reg. No. 38,080

"Express Mail" mailing label number EL874326426US

Date of Deposit November 26, 2001

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" services under 37 C.F.R. 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Typed Name of Person Mailing Paper or Fee: Christine Docherty

Signature: 

15/pts

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

5

The Field of the Invention

The present invention generally relates to memory management in computer systems, and more particularly to allocation of memory in a distributed computer system including local memory access protection and remote memory access protection.

10

Background of the Invention

Conventional computer systems typically employ region-based memory management techniques. In the region-based memory management technique, a large portion of memory (e.g., one megabyte) is allocated to an application. This large allocated memory portion is referred to as the application's memory domain. Conventional computer systems allocate the large portion of memory for the application by trapping to an operating system kernel process. The trap to the operating system kernel process is performed to change memory access permissions and the like.

20

Conventional computer systems employ protection tables storing information related to memory page access rights, and page tables storing information for mapping virtual to physical addresses.

These conventional memory management techniques are cumbersome when multiple applications share memory and require access to the same shared memory segments. For example, when one application accesses a shared memory segment, and then another application uses the same shared memory segment, the virtual to physical mappings in the page table change. These virtual to physical mapping changes are difficult to track by the multiple applications sharing the same memory segment.

30

For reasons stated above and for other reasons presented in greater detail in the Description of the Preferred Embodiment section of the present

specification, there is a need for an improved memory management technique for permitting memory access protection at a lower level granularity than currently employed by conventional memory management techniques which always allocate a large portion of memory to an application. In addition, an
5 improved memory management technique is needed which efficiently supports shared memory. Moreover, the improved memory management technique should permit virtual addresses to be used by applications in a shared memory system so that the applications do not need to deal with physical addresses. Furthermore, an improved memory management technique is desired which
10 permits some types of memory allocation without context switching and without trapping to an operating system kernel process.

Summary of the Invention

The present invention provides a method of managing memory in a distributed computer system. The method comprises binding a remote key to a
5 first address representing a contiguous memory address range at a first endnode. The bound remote key and first address are sent from the first endnode to a second endnode on a communication fabric. A remote direct memory access operation is performed from the second endnode to access the contiguous memory address range including sending the bound remote key and the first
10 address from the second endnode to the first endnode on the communication fabric.

Brief Description of the Drawings

Figure 1 is a diagram of a distributed computer system for implementing the present invention.

5 Figure 2 is a diagram of an example host processor node for the computer system of Figure 1.

Figure 3 is a diagram of a portion of a distributed computer system employing a reliable connection service to communicate between distributed processes.

10 Figure 4 is a diagram of a portion of distributed computer system employing a reliable datagram service to communicate between distributed processes.

Figure 5 is a diagram of an example host processor node for operation in a distributed computer system implementing the present invention.

15 Figure 6 is a diagram of a portion of a distributed computer system illustrating subnets in the distributed computer system.

Figure 7 is a diagram of a switch for use in a distributed computer system implemented the present invention.

Figure 8 is a diagram of a portion of a distributed computer system.

20 Figure 9A is a diagram of a work queue element (WQE) for operation in the distributed computer system of Figure 8.

Figure 9B is a diagram of the packetization process of a message created by the WQE of Figure 9A into frames and flits.

25 Figure 10A is a diagram of a message being transmitted with a reliable transport service illustrating frame transactions.

Figure 10B is a diagram illustrating a reliable transport service illustrating flit transactions associated with the frame transactions of Figure 10A.

Figure 11 is a diagram of a layered architecture for implementing the present invention.

30 Figure 12 is a diagram of a computer system for illustrating memory access protection mechanisms.

Figure 13 is a diagram of an example effective address.

Figure 14 is a diagram of an example segment table entry.

Figure 15 is a diagram of an example virtual address.

Figure 16 is a diagram of an example page table entry.

5 Figure 17 is flow diagram of a mechanism for managing local memory access protection.

Figure 18 is a diagram of a distributed computer system having an embodiment of a mechanism for managing remote memory access protection according to the present invention.

10

Description of the Preferred Embodiments

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural or logical changes may be made without departing from the scope of the present invention. The following detailed description, therefore, is not to be taken in a limiting sense, and the scope of the present invention is defined by the appended claims.

15

20

One embodiment of the present invention is directed to a memory management technique employing remote keys to allocate a small region of memory referred to as a memory window without trapping to an operating system kernel process. In such a memory management technique according to the present invention, a range of memory in the memory window or the entire memory window can be allocated by binding a remote key to the range of memory in the memory window or the entire memory window to permit memory access protection at a byte level granularity. One embodiment of a memory management technique according to the present invention provides an extremely light-weight technique for supporting shared memory. One embodiment of a memory management technique according to the present invention permits

25

30

remote memory management where user applications only deal with virtual addresses.

An example embodiment of a distributed computer system is illustrated generally at 30 in Figure 1. Distributed computer system 30 is provided merely for illustrative purposes, and the embodiments of the present invention described below can be implemented on computer systems of numerous other types and configurations. For example, computer systems implementing the present invention can range from a small server with one processor and a few input/output (I/O) adapters to massively parallel supercomputer systems with hundreds or thousands of processors and thousands of I/O adapters. Furthermore, the present invention can be implemented in an infrastructure of remote computer systems connected by an internet or intranet.

Distributed computer system 30 includes a system area network (SAN) 32 which is a high-bandwidth, low-latency network interconnecting nodes within distributed computer system 30. A node is herein defined to be any device attached to one or more links of a network and forming the origin and/or destination of messages within the network. In the example distributed computer system 30, nodes include host processors 34a-34d; redundant array independent disk (RAID) subsystem 33; and I/O adapters 35a and 35b. The nodes illustrated in Figure 1 are for illustrative purposes only, as SAN 32 can connect any number and any type of independent processor nodes, I/O adapter nodes, and I/O device nodes. Any one of the nodes can function as an endnode, which is herein defined to be a device that originates or finally consumes messages or frames in the distributed computer system.

A message is herein defined to be an application-defined unit of data exchange, which is a primitive unit of communication between cooperating sequential processes. A frame is herein defined to be one unit of data encapsulated by a physical network protocol header and/or trailer. The header generally provides control and routing information for directing the frame through SAN 32. The trailer generally contains control and cyclic redundancy check (CRC) data for ensuring packets are not delivered with corrupted contents.

SAN 32 is the communications and management infrastructure supporting both I/O and interprocess communication (IPC) within distributed computer system 30. SAN 32 includes a switched communications fabric (SAN FABRIC) allowing many devices to concurrently transfer data with high-
5 bandwidth and low latency in a secure, remotely managed environment. Endnodes can communicate over multiple ports and utilize multiple paths through the SAN fabric. The multiple ports and paths through SAN 32 can be employed for fault tolerance and increased bandwidth data transfers.

SAN 32 includes switches 36 and routers 38. A switch is herein defined
10 to be a device that connects multiple links 40 together and allows routing of frames from one link 40 to another link 40 within a subnet using a small header destination ID field. A router is herein defined to be a device that connects multiple links 40 together and is capable of routing frames from one link 40 in a first subnet to another link 40 in a second subnet using a large header destination
15 address or source address.

In one embodiment, a link 40 is a full duplex channel between any two network fabric elements, such as endnodes, switches 36, or routers 38. Example suitable links 40 include, but are not limited to, copper cables, optical cables, and printed circuit copper traces on backplanes and printed circuit boards.

20 Endnodes, such as host processor endnodes 34 and I/O adapter endnodes 35, generate request frames and return acknowledgment frames. By contrast, switches 36 and routers 38 do not generate and consume frames. Switches 36 and routers 38 simply pass frames along. In the case of switches 36, the frames are passed along unmodified. For routers 38, the network header is modified
25 slightly when the frame is routed. Endnodes, switches 36, and routers 38 are collectively referred to as end stations.

In distributed computer system 30, host processor nodes 34a-34d and RAID subsystem node 33 include at least one system area network interface controller (SANIC) 42. In one embodiment, each SANIC 42 is an endpoint that
30 implements the SAN 32 interface in sufficient detail to source or sink frames transmitted on the SAN fabric. The SANICs 42 provide an interface to the host

processors and I/O devices. In one embodiment the SANIC is implemented in hardware. In this SANIC hardware implementation, the SANIC hardware offloads much of CPU and I/O adapter communication overhead. This hardware implementation of the SANIC also permits multiple concurrent communications
5 over a switched network without the traditional overhead associated with communicating protocols. In one embodiment, SAN 32 provides the I/O and IPC clients of distributed computer system 30 zero processor-copy data transfers without involving the operating system kernel process, and employs hardware to provide reliable, fault tolerant communications.

10 As indicated in Figure 1, router 38 is coupled to wide area network (WAN) and/or local area network (LAN) connections to other hosts or other routers 38.

The host processors 34a-34d include central processing units (CPUs) 44 and memory 46.

15 I/O adapters 35a and 35b include an I/O adapter backplane 48 and multiple I/O adapter cards 50. Example adapter cards 50 illustrated in Figure 1 include an SCSI adapter card; an adapter card to fiber channel hub and FC-AL devices; an Ethernet adapter card; and a graphics adapter card. Any known type of adapter card can be implemented. I/O adapters 35a and 35b also include a
20 switch 36 in the I/O adapter backplane 48 to couple the adapter cards 50 to the SAN 32 fabric.

RAID subsystem 33 includes a microprocessor 52, memory 54, read/write circuitry 56, and multiple redundant storage disks 58.

SAN 32 handles data communications for I/O and IPC in distributed
25 computer system 30. SAN 32 supports high-bandwidth and scalability required for I/O and also supports the extremely low latency and low CPU overhead required for IPC. User clients can bypass the operating system kernel process and directly access network communication hardware, such as SANICs 42 which enable efficient message passing protocols. SAN 32 is suited to current
30 computing models and is a building block for new forms of I/O and computer cluster communication. SAN 32 allows I/O adapter nodes to communicate

among themselves or communicate with any or all of the processor nodes in distributed computer system 30. With an I/O adapter attached to SAN 32, the resulting I/O adapter node has substantially the same communication capability as any processor node in distributed computer system 30.

5

Channel and Memory Semantics

In one embodiment, SAN 32 supports channel semantics and memory semantics. Channel semantics is sometimes referred to as send/receive or push communication operations, and is the type of communications employed in a traditional I/O channel where a source device pushes data and a destination device determines the final destination of the data. In channel semantics, the frame transmitted from a source process specifies a destination processes' communication port, but does not specify where in the destination processes' memory space the frame will be written. Thus, in channel semantics, the destination process pre-allocates where to place the transmitted data.

In memory semantics, a source process directly reads or writes the virtual address space of a remote node destination process. The remote destination process need only communicate the location of a buffer for data, and does not need to be involved with the transfer of any data. Thus, in memory semantics, a source process sends a data frame containing the destination buffer memory address of the destination process. In memory semantics, the destination process previously grants permission for the source process to access its memory.

Channel semantics and memory semantics are typically both necessary for I/O and IPC. A typical I/O operation employs a combination of channel and memory semantics. In an illustrative example I/O operation of distributed computer system 30, host processor 34a initiates an I/O operation by using channel semantics to send a disk write command to I/O adapter 35b. I/O adapter 35b examines the command and uses memory semantics to read the data buffer directly from the memory space of host processor 34a. After the data buffer is read, I/O adapter 35b employs channel semantics to push an I/O completion message back to host processor 34a.

In one embodiment, distributed computer system 30 performs operations that employ virtual addresses and virtual memory protection mechanisms to ensure correct and proper access to all memory. In one embodiment, applications running in distributed computer system 30 are not required to use physical addressing for any operations.

Queue Pairs

An example host processor node 34 is generally illustrated in Figure 2. Host processor node 34 includes a process A indicated at 60 and a process B indicated at 62. Host processor node 34 includes SANIC 42. Host processor node 34 also includes queue pairs (QP's) 64a and 64b which provide communication between process 60 and SANIC 42. Host processor node 34 also includes QP 64c which provides communication between process 62 and SANIC 42. A single SANIC, such as SANIC 42 in a host processor 34, can support thousands of QPs. By contrast, a SAN interface in an I/O adapter 35 typically supports less than ten QPs.

Each QP 64 includes a send work queue 66 and a receive work queue 68. A process, such as processes 60 and 62, calls an operating-system specific programming interface which is herein referred to as verbs, which place work items, referred to as work queue elements (WQEs) onto a QP 64. A WQE is executed by hardware in SANIC 42. SANIC 42 is coupled to SAN 32 via physical link 40. Send work queue 66 contains WQEs that describe data to be transmitted on the SAN 32 fabric. Receive work queue 68 contains WQEs that describe where to place incoming data from the SAN 32 fabric.

Host processor node 34 also includes completion queue 70a interfacing with process 60 and completion queue 70b interfacing with process 62. The completion queues 70 contain information about completed WQEs. The completion queues are employed to create a single point of completion notification for multiple QPs. A completion queue entry is a data structure on a completion queue 70 that describes a completed WQE. The completion queue entry contains sufficient information to determine the QP that holds the

completed WQE. A completion queue context is a block of information that contains pointers to, length, and other information needed to manage the individual completion queues.

Example WQEs include work items that initiate data communications employing channel semantics or memory semantics; work items that are instructions to hardware in SANIC 42 to set or alter remote memory access protections; and work items to delay the execution of subsequent WQEs posted in the same send work queue 66.

More specifically, example WQEs supported for send work queues 66 are as follows. A send buffer WQE is a channel semantic operation to push a local buffer to a remote QP's receive buffer. The send buffer WQE includes a gather list to combine several virtual contiguous local buffers into a single message that is pushed to a remote QP's receive buffer. The local buffer virtual addresses are in the address space of the process that created the local QP.

A remote direct memory access (RDMA) read WQE provides a memory semantic operation to read a virtually contiguous buffer on a remote node. The RDMA read WQE reads a virtually contiguous buffer on a remote endnode and writes the data to a scatter list of virtually contiguous local memory buffers. Similar to the send buffer WQE, the local buffer for the RDMA read WQE is in the address space of the process that created the local QP. The remote buffer is in the virtual address space of the process owning the remote QP targeted by the RDMA read WQE.

A RDMA write WQE provides a memory semantic operation to write a virtually contiguous buffer on a remote node. The RDMA write WQE contains a scatter list of locally virtually contiguous buffers and the virtual address of the remote buffer into which the local buffers are written.

A RDMA FetchOp WQE provides a memory semantic operation to perform an atomic operation on a remote word. The RDMA FetchOp WQE is a combined RDMA read, modify, and RDMA write operation. The RDMA FetchOp WQE can support several read-modify-write operations, such as Compare and Swap if equal.

A bind/unbind remote access key (RKey) WQE provides a command to SANIC hardware to modify the association of a RKey with a local virtually contiguous buffer. The RKey is part of each RDMA access and is used to validate that the remote process has permitted access to the buffer.

5 A delay WQE provides a command to SANIC hardware to delay processing of the QP's WQEs for a specific time interval. The delay WQE permits a process to meter the flow of operations into the SAN fabric.

10 In one embodiment, receive queues 68 only support one type of WQE, which is referred to as a receive buffer WQE. The receive buffer WQE provides a channel semantic operation describing a local buffer into which incoming send messages are written. The receive buffer WQE includes a scatter list describing several virtually contiguous local buffers. An incoming send message is written to these buffers. The buffer virtual addresses are in the address space of the process that created the local QP.

15 For IPC, a user-mode software process transfers data through QPs 64 directly from where the buffer resides in memory. In one embodiment, the transfer through the QPs bypasses the operating system and consumes few host instruction cycles. QPs 64 permit zero processor-copy data transfer with no operating system kernel involvement. The zero processor-copy data transfer
20 provides for efficient support of high-bandwidth and low-latency communication.

Transport Services

25 When a QP 64 is created, the QP is set to provide a selected type of transport service. In one embodiment, a distributed computer system implementing the present invention supports four types of transport services.

30 A portion of a distributed computer system employing a reliable connection service to communicate between distributed processes is illustrated generally at 100 in Figure 3. Distributed computer system 100 includes a host processor node 102, a host processor node 104, and a host processor node 106. Host processor node 102 includes a process A indicated at 108. Host processor

node 104 includes a process B indicated at 110 and a process C indicated at 112.
Host processor node 106 includes a process D indicated at 114.

Host processor node 102 includes a QP 116 having a send work queue 116a and a receive work queue 116b; a QP 118 having a send work queue 118a
5 and receive work queue 118b; and a QP 120 having a send work queue 120a and a receive work queue 120b which facilitate communication to and from process A indicated at 108. Host processor node 104 includes a QP 122 having a send work queue 122a and receive work queue 122b for facilitating communication to and from process B indicated at 110. Host processor node 104 includes a QP
10 124 having a send work queue 124a and receive work queue 124b for facilitating communication to and from process C indicated at 112. Host processor node 106 includes a QP 126 having a send work queue 126a and receive work queue 126b for facilitating communication to and from process D indicated at 114.

The reliable connection service of distributed computer system 100
15 associates a local QP with one and only one remote QP. Thus, QP 116 is connected to QP 122 via a non-sharable resource connection 128 having a non-sharable resource connection 128a from send work queue 116a to receive work queue 122b and a non-sharable resource connection 128b from send work queue 122a to receive work queue 116b. QP 118 is connected to QP 124 via a non-
20 sharable resource connection 130 having a non-sharable resource connection 130a from send work queue 118a to receive work queue 124b and a non-sharable resource connection 130b from send work queue 124a to receive work queue 118b. QP 120 is connected to QP 126 via a non-sharable resource connection 132 having a non-sharable resource connection 132a from send work queue 120a
25 to receive work queue 126b and a non-sharable resource connection 132b from send work queue 126a to receive work queue 120b.

A send buffer WQE placed on one QP in a reliable connection service causes data to be written into the receive buffer of the connected QP. RDMA operations operate on the address space of the connected QP.

30 The reliable connection service requires a process to create a QP for each process which is to communicate with over the SAN fabric. Thus, if each of N

host processor nodes contain M processes, and all M processes on each node wish to communicate with all the processes on all the other nodes, each host processor node requires $M^2 \times (N - 1)$ QPs. Moreover, a process can connect a QP to another QP on the same SANIC.

5 In one embodiment, the reliable connection service is made reliable because hardware maintains sequence numbers and acknowledges all frame transfers. A combination of hardware and SAN driver software retries any failed communications. The process client of the QP obtains reliable communications even in the presence of bit errors, receive buffer underruns, and network
10 congestion. If alternative paths exist in the SAN fabric, reliable communications can be maintained even in the presence of failures of fabric switches or links.

 In one embodiment, acknowledgements are employed to deliver data reliably across the SAN fabric. In one embodiment, the acknowledgement is not a process level acknowledgment, because the acknowledgment does not validate
15 the receiving process has consumed the data. Rather, the acknowledgment only indicates that the data has reached its destination.

 A portion of a distributed computer system employing a reliable datagram service to communicate between distributed processes is illustrated generally at 150 in Figure 4. Distributed computer system 150 includes a host
20 processor node 152, a host processor node 154, and a host processor node 156. Host processor node 152 includes a process A indicated at 158. Host processor node 154 includes a process B indicated at 160 and a process C indicated at 162. Host processor node 156 includes a process D indicated at 164.

 Host processor node 152 includes QP 166 having send work queue 166a
25 and receive work queue 166b for facilitating communication to and from process A indicated at 158. Host processor node 154 includes QP 168 having send work queue 168a and receive work queue 168b for facilitating communication from and to process B indicated at 160. Host processor node 154 includes QP 170 having send work queue 170a and receive work queue 170b for facilitating
30 communication from and to process C indicated at 162. Host processor node 156 includes QP 172 having send work queue 172a and receive work queue

172b for facilitating communication from and to process D indicated at 164. In the reliable datagram service implemented in distributed computer system 150, the QPs are coupled in what is referred to as a connectionless transport service.

For example, a reliable datagram service 174 couples QP 166 to QPs 168, 170, and 172. Specifically, reliable datagram service 174 couples send work queue 166a to receive work queues 168b, 170b, and 172b. Reliable datagram service 174 also couples send work queues 168a, 170a, and 172a to receive work queue 166b.

The reliable datagram service permits a client process of one QP to communicate with any other QP on any other remote node. At a receive work queue, the reliable datagram service permits incoming messages from any send work queue on any other remote node.

In one embodiment, the reliable datagram service employs sequence numbers and acknowledgments associated with each message frame to ensure the same degree of reliability as the reliable connection service. End-to-end (EE) contexts maintain end-to-end specific state to keep track of sequence numbers, acknowledgments, and time-out values. The end-to-end state held in the EE contexts is shared by all the connectionless QPs communicating between a pair of endnodes. Each endnode requires at least one EE context for every endnode it wishes to communicate with in the reliable datagram service (e.g., a given endnode requires at least N EE contexts to be able to have reliable datagram service with N other endnodes).

The reliable datagram service greatly improves scalability because the reliable datagram service is connectionless. Therefore, an endnode with a fixed number of QPs can communicate with far more processes and endnodes with a reliable datagram service than with a reliable connection transport service. For example, if each of N host processor nodes contain M processes, and all M processes on each node wish to communicate with all the processes on all the other nodes, the reliable connection service requires $M^2 \times (N - 1)$ QPs on each node. By comparison, the connectionless reliable datagram service only requires

M QPs + (N - 1) EE contexts on each node for exactly the same communications.

A third type of transport service for providing communications is a unreliable datagram service. Similar to the reliable datagram service, the
5 unreliable datagram service is connectionless. The unreliable datagram service is employed by management applications to discover and integrate new switches, routers, and endnodes into a given distributed computer system. The unreliable datagram service does not provide the reliability guarantees of the reliable connection service and the reliable datagram service. The unreliable datagram
10 service accordingly operates with less state information maintained at each endnode.

A fourth type of transport service is referred to as raw datagram service and is technically not a transport service. The raw datagram service permits a QP to send and to receive raw datagram frames. The raw datagram mode of
15 operation of a QP is entirely controlled by software. The raw datagram mode of the QP is primarily intended to allow easy interfacing with traditional internet protocol, version 6 (IPv6) LAN-WAN networks, and further allows the SANIC to be used with full software protocol stacks to access transmission control protocol (TCP), user datagram protocol (UDP), and other standard
20 communication protocols. Essentially, in the raw datagram service, SANIC hardware generates and consumes standard protocols layered on top of IPv6, such as TCP and UDP. The frame header can be mapped directly to and from an IPv6 header. Native IPv6 frames can be bridged into the SAN fabric and delivered directly to a QP to allow a client process to support any transport
25 protocol running on top of IPv6. A client process can register with SANIC hardware in order to direct datagrams for a particular upper level protocol (e.g., TCP and UDP) to a particular QP. SANIC hardware can demultiplex incoming IPv6 streams of datagrams based on a next header field as well as the destination IP address.

30

SANIC and I/O Adapter Endnodes

20140102608660

An example host processor node is generally illustrated at 200 in Figure 5. Host processor node 200 includes a process A indicated at 202, a process B indicated at 204, and a process C indicated at 206. Host processor 200 includes a SANIC 208 and a SANIC 210. As discussed above, a host processor endnode or an I/O adapter endnode can have one or more SANICs. SANIC 208 includes a SAN link level engine (LLE) 216 for communicating with SAN fabric 224 via link 217 and an LLE 218 for communicating with SAN fabric 224 via link 219. SANIC 210 includes an LLE 220 for communicating with SAN fabric 224 via link 221 and an LLE 222 for communicating with SAN fabric 224 via link 223. SANIC 208 communicates with process A indicated at 202 via QPs 212a and 212b. SANIC 208 communicates with process B indicated at 204 via QPs 212c-212n. Thus, SANIC 208 includes N QPs for communicating with processes A and B. SANIC 210 includes QPs 214a and 214b for communicating with process B indicated at 204. SANIC 210 includes QPs 214c-214n for communicating with process C indicated at 206. Thus, SANIC 210 includes N QPs for communicating with processes B and C.

An LLE runs link level protocols to couple a given SANIC to the SAN fabric. RDMA traffic generated by a SANIC can simultaneously employ multiple LLEs within the SANIC which permits striping across LLEs. Striping refers to the dynamic sending of frames within a single message to an endnode's QP through multiple fabric paths. Striping across LLEs increases the bandwidth for a single QP as well as provides multiple fault tolerant paths. Striping also decreases the latency for message transfers. In one embodiment, multiple LLEs in a SANIC are not visible to the client process generating message requests. When a host processor includes multiple SANICs, the client process must explicitly move data on the two SANICs in order to gain parallelism. A single QP cannot be shared by SANICs. Instead a QP is owned by one local SANIC.

The following is an example naming scheme for naming and identifying endnodes in one embodiment of a distributed computer system according to the present invention. A host name provides a logical identification for a host node, such as a host processor node or I/O adapter node. The host name identifies the

endpoint for messages such that messages are destined for processes residing on an endnode specified by the host name. Thus, there is one host name per node, but a node can have multiple SANICs.

A globally unique ID (GUID) identifies a transport endpoint. A transport
5 endpoint is the device supporting the transport QPs. There is one GUID associated with each SANIC.

A local ID refers to a short address ID used to identify a SANIC within a single subnet. In one example embodiment, a subnet has up to 2^{16} endnodes, switches, and routers, and the local ID (LID) is accordingly 16 bits. A source
10 LID (SLID) and a destination LID (DLID) are the source and destination LIDs used in a local network header. A LLE has a single LID associated with the LLE, and the LID is only unique within a given subnet. One or more LIDs can be associated with each SANIC.

An internet protocol (IP) address (e.g., a 128 bit IPv6 ID) addresses a
15 SANIC. The SANIC, however, can have one or more IP addresses associated with the SANIC. The IP address is used in the global network header when routing frames outside of a given subnet. LIDs and IP addresses are network endpoints and are the target of frames routed through the SAN fabric. All IP addresses (e.g., IPv6 addresses) within a subnet share a common set of high
20 order address bits.

In one embodiment, the LLE is not named and is not architecturally visible to a client process. In this embodiment, management software refers to LLEs as an enumerated subset of the SANIC.

25 Switches and Routers

A portion of a distributed computer system is generally illustrated at 250 in Figure 6. Distributed computer system 250 includes a subnet A indicated at 252 and a subnet B indicated at 254. Subnet A indicated at 252 includes a host processor node 256 and a host processor node 258. Subnet B indicated at 254
30 includes a host processor node 260 and host processor node 262. Subnet A indicated at 252 includes switches 264a-264c. Subnet B indicated at 254

includes switches 266a-266c. Each subnet within distributed computer system 250 is connected to other subnets with routers. For example, subnet A indicated at 252 includes routers 268a and 268b which are coupled to routers 270a and 270b of subnet B indicated at 254. In one example embodiment, a subnet has up
5 to 2^{16} endnodes, switches, and routers.

A subnet is defined as a group of endnodes and cascaded switches that is managed as a single unit. Typically, a subnet occupies a single geographic or functional area. For example, a single computer system in one room could be defined as a subnet. In one embodiment, the switches in a subnet can perform
10 very fast worm-hole or cut-through routing for messages.

A switch within a subnet examines the DLID that is unique within the subnet to permit the switch to quickly and efficiently route incoming message frames. In one embodiment, the switch is a relatively simple circuit, and is typically implemented as a single integrated circuit. A subnet can have hundreds
15 to thousands of endnodes formed by cascaded switches.

As illustrated in Figure 6, for expansion to much larger systems, subnets are connected with routers, such as routers 268 and 270. The router interprets the IP destination ID (e.g., IPv6 destination ID) and routes the IP like frame.

In one embodiment, switches and routers degrade when links are over
20 utilized. In this embodiment, link level back pressure is used to temporarily slow the flow of data when multiple input frames compete for a common output. However, link or buffer contention does not cause loss of data. In one embodiment, switches, routers, and endnodes employ a link protocol to transfer data. In one embodiment, the link protocol supports an automatic error retry. In
25 this example embodiment, link level acknowledgments detect errors and force retransmission of any data impacted by bit errors. Link-level error recovery greatly reduces the number of data errors that are handled by the end-to-end protocols. In one embodiment, the user client process is not involved with error recovery no matter if the error is detected and corrected by the link level
30 protocol or the end-to-end protocol.

An example embodiment of a switch is generally illustrated at 280 in Figure 7. Each I/O path on a switch or router has an LLE. For example, switch 280 includes LLEs 282a-282h for communicating respectively with links 284a-284h.

5 The naming scheme for switches and routers is similar to the above-described naming scheme for endnodes. The following is an example switch and router naming scheme for identifying switches and routers in the SAN fabric. A switch name identifies each switch or group of switches packaged and managed together. Thus, there is a single switch name for each switch or group of
10 switches packaged and managed together.

Each switch or router element has a single unique GUID. Each switch has one or more LIDs and IP addresses (e.g., IPv6 addresses) that are used as an endnode for management frames.

Each LLE is not given an explicit external name in the switch or router.
15 Since links are point-to-point, the other end of the link does not need to address the LLE.

Virtual Lanes

Switches and routers employ multiple virtual lanes within a single
20 physical link. As illustrated in Figure 6, physical links 272 connect endnodes, switches, and routers within a subnet. WAN or LAN connections 274 typically couple routers between subnets. Frames injected into the SAN fabric follow a particular virtual lane from the frame's source to the frame's destination. At any one time, only one virtual lane makes progress on a given physical link. Virtual
25 lanes provide a technique for applying link level flow control to one virtual lane without affecting the other virtual lanes. When a frame on one virtual lane blocks due to contention, quality of service (QoS), or other considerations, a frame on a different virtual lane is allowed to make progress.

Virtual lanes are employed for numerous reasons, some of which are as
30 follows. Virtual lanes provide QoS. In one example embodiment, certain virtual lanes are reserved for high priority or isochronous traffic to provide QoS.

Virtual lanes provide deadlock avoidance. Virtual lanes allow topologies that contain loops to send frames across all physical links and still be assured the loops won't cause back pressure dependencies that might result in deadlock.

Virtual lanes alleviate head-of-line blocking. With virtual lanes, a
5 blocked frames can pass a temporarily stalled frame that is destined for a different final destination.

In one embodiment, each switch includes its own crossbar switch. In this embodiment, a switch propagates data from only one frame at a time, per virtual lane through its crossbar switch. In another words, on any one virtual lane, a
10 switch propagates a single frame from start to finish. Thus, in this embodiment, frames are not multiplexed together on a single virtual lane.

Paths in SAN fabric

Referring to Figure 6, within a subnet, such as subnet A indicated at 252
15 or subnet B indicated at 254, a path from a source port to a destination port is determined by the LID of the destination SANIC port. Between subnets, a path is determined by the IP address (e.g., IPv6 address) of the destination SANIC port.

In one embodiment, the paths used by the request frame and the request
20 frame's corresponding positive acknowledgment (ACK) or negative acknowledgment (NAK) frame are not required to be symmetric. In one embodiment employing oblivious routing, switches select an output port based on the DLID. In one embodiment, a switch uses one set of routing decision criteria for all its input ports. In one example embodiment, the routing decision
25 criteria is contained in one routing table. In an alternative embodiment, a switch employs a separate set of criteria for each input port.

Each port on an endnode can have multiple IP addresses. Multiple IP addresses can be used for several reasons, some of which are provided by the following examples. In one embodiment, different IP addresses identify
30 different partitions or services on an endnode. In one embodiment, different IP

addresses are used to specify different QoS attributes. In one embodiment, different IP addresses identify different paths through intra-subnet routes.

In one embodiment, each port on an endnode can have multiple LIDs. Multiple LIDs can be used for several reasons some of which are provided by
5 the following examples. In one embodiment, different LIDs identify different partitions or services on an endnode. In one embodiment, different LIDs are used to specify different QoS attributes. In one embodiment, different LIDs specify different paths through the subnet.

A one-to-one correspondence does not necessarily exist between LIDs
10 and IP addresses, because a SANIC can have more or less LIDs than IP addresses for each port. For SANICs with redundant ports and redundant conductivity to multiple SAN fabrics, SANICs can, but are not required to, use the same LID and IP address on each of its ports.

15 Data Transactions

Referring to Figure 1, a data transaction in distributed computer system
30 is typically composed of several hardware and software steps. A client process of a data transport service can be a user-mode or a kernel-mode process. The client process accesses SANIC 42 hardware through one or more QPs, such
20 as QPs 64 illustrated in Figure 2. The client process calls an operating-system specific programming interface which is herein referred to as verbs. The software code implementing the verbs internally posts a WQE to the given QP work queue.

There are many possible methods of posting a WQE and there are many
25 possible WQE formats, which allow for various cost/performance design points, but which do not affect interoperability. A user process, however, must communicate to verbs in a well-defined manner, and the format and protocols of data transmitted across the SAN fabric must be sufficiently specified to allow devices to interoperate in a heterogeneous vendor environment.

30 In one embodiment, SANIC hardware detects WQE posting and accesses the WQE. In this embodiment, the SANIC hardware translates and validates the

WQEs virtual addresses and accesses the data. In one embodiment, an outgoing message buffer is split into one or more frames. In one embodiment, the SANIC hardware adds a transport header and a network header to each frame. The transport header includes sequence numbers and other transport information.

- 5 The network header includes the destination IP address or the DLID or other suitable destination address information. The appropriate local or global network header is added to a given frame depending on if the destination endnode resides on the local subnet or on a remote subnet.

A frame is a unit of information that is routed through the SAN fabric.

- 10 The frame is an endnode-to-endnode construct, and is thus created and consumed by endnodes. Switches and routers neither generate nor consume request frames or acknowledgment frames. Instead switches and routers simply move request frames or acknowledgment frames closer to the ultimate destination. Routers, however, modify the frame's network header when the
- 15 frame crosses a subnet boundary. In traversing a subnet, a single frame stays on a single virtual lane.

- When a frame is placed onto a link, the frame is further broken down into flits. A flit is herein defined to be a unit of link-level flow control and is a unit of transfer employed only on a point-to-point link. The flow of flits is subject to
- 20 the link-level protocol which can perform flow control or retransmission after an error. Thus, flit is a link-level construct that is created at each endnode, switch, or router output port and consumed at each input port. In one embodiment, a flit contains a header with virtual lane error checking information, size information, and reverse channel credit information.

- 25 If a reliable transport service is employed, after a request frame reaches its destination endnode, the destination endnode sends an acknowledgment frame back to the sender endnode. The acknowledgment frame permits the requestor to validate that the request frame reached the destination endnode. An acknowledgment frame is sent back to the requestor after each request frame.

- 30 The requestor can have multiple outstanding requests before it receives any

acknowledgments. In one embodiment, the number of multiple outstanding requests is determined when a QP is created.

Example Request and Acknowledgment Transactions

5 Figures 8, 9A, 9B, 10A, and 10B together illustrate example request and acknowledgment transactions. In Figure 8, a portion of a distributed computer system is generally illustrated at 300. Distributed computer system 300 includes a host processor node 302 and a host processor node 304. Host processor node 302 includes a SANIC 306. Host processor node 304 includes a SANIC 308.
10 Distributed computer system 300 includes a SAN fabric 309 which includes a switch 310 and a switch 312. SAN fabric 309 includes a link 314 coupling SANIC 306 to switch 310; a link 316 coupling switch 310 to switch 312; and a link 318 coupling SANIC 308 to switch 312.

 In the example transactions, host processor node 302 includes a client
15 process A indicated at 320. Host processor node 304 includes a client process B indicated at 322. Client process 320 interacts with SANIC hardware 306 through QP 324. Client process 322 interacts with SANIC hardware 308 through QP 326. QP 324 and 326 are software data structures. QP 324 includes send work queue 324a and receive work queue 324b. QP 326 includes send
20 work queue 326a and receive work queue 326b.

 Process 320 initiates a message request by posting WQEs to send queue 324a. Such a WQE is illustrated at 330 in Figure 9A. The message request of client process 320 is referenced by a gather list 332 contained in send WQE 330. Each entry in gather list 332 points to a virtually contiguous buffer in the local
25 memory space containing a part of the message, such as indicated by virtual contiguous buffers 334a-334d, which respectively hold message 0, parts 0, 1, 2, and 3.

 Referring to Figure 9B, hardware in SANIC 306 reads WQE 330 and packetizes the message stored in virtual contiguous buffers 334a-334d into
30 frames and flits. As illustrated in Figure 9B, all of message 0, part 0 and a portion of message 0, part 1 are packetized into frame 0, indicated at 336a. The

rest of message 0, part 1 and all of message 0, part 2, and all of message 0, part 3 are packetized into frame 1, indicated at 336b. Frame 0 indicated at 336a includes network header 338a and transport header 340a. Frame 1 indicated at 336b includes network header 338b and transport header 340b.

5 As indicated in Figure 9B, frame 0 indicated at 336a is partitioned into flits 0-3, indicated respectively at 342a-342d. Frame 1 indicated at 336b is partitioned into flits 4-7 indicated respectively at 342e - 342h. Flits 342a through 342h respectively include flit headers 344a-344h.

10 Frames are routed through the SAN fabric, and for reliable transfer services, are acknowledged by the final destination endnode. If not successively acknowledged, the frame is retransmitted by the source endnode. Frames are generated by source endnodes and consumed by destination endnodes. The switches and routers in the SAN fabric neither generate nor consume frames.

15 Flits are the smallest unit of flow control in the network. Flits are generated and consumed at each end of a physical link. Flits are acknowledged at the receiving end of each link and are retransmitted in response to an error.

Referring to Figure 10A, the send request message 0 is transmitted from SANIC 306 in host processor node 302 to SANIC 308 in host processor node 304 as frames 0 indicated at 336a and frame 1 indicated at 336b. ACK frames 20 346a and 346b, corresponding respectively to request frames 336a and 336b, are transmitted from SANIC 308 in host processor node 304 to SANIC 306 in host processor node 302.

In Figure 10A, message 0 is being transmitted with a reliable transport service. Each request frame is individually acknowledged by the destination 25 endnode (e.g., SANIC 308 in host processor node 304).

Figure 10B illustrates the flits associated with the request frames 336 and acknowledgment frames 346 illustrated in Figure 10A passing between the host processor endnodes 302 and 304 and the switches 310 and 312. As illustrated in Figure 10B, an ACK frame fits inside one flit. In one embodiment, one 30 acknowledgment flit acknowledges several flits.

As illustrated in Figure 10B, flits 342a-h are transmitted from SANIC 306 to switch 310. Switch 310 consumes flits 342a-h at its input port, creates flits 348a-h at its output port corresponding to flits 342a-h, and transmits flits 348a-h to switch 312. Switch 312 consumes flits 348a-h at its input port, creates
5 flits 350a-h at its output port corresponding to flits 348a-h, and transmits flits 350a-h to SANIC 308. SANIC 308 consumes flits 350a-h at its input port. An acknowledgment flit is transmitted from switch 310 to SANIC 306 to acknowledge the receipt of flits 342a-h. An acknowledgment flit 354 is transmitted from switch 312 to switch 310 to acknowledge the receipt of flits
10 348a-h. An acknowledgment flit 356 is transmitted from SANIC 308 to switch 312 to acknowledge the receipt of flits 350a-h.

Acknowledgment frame 346a fits inside of flit 358 which is transmitted from SANIC 308 to switch 312. Switch 312 consumes flits 358 at its input port, creates flit 360 corresponding to flit 358 at its output port, and transmits flit 360
15 to switch 310. Switch 310 consumes flit 360 at its input port, creates flit 362 corresponding to flit 360 at its output port, and transmits flit 362 to SANIC 306. SANIC 306 consumes flit 362 at its input port. Similarly, SANIC 308 transmits acknowledgment frame 346b in flit 364 to switch 312. Switch 312 creates flit 366 corresponding to flit 364, and transmits flit 366 to switch 310. Switch 310
20 creates flit 368 corresponding to flit 366, and transmits flit 368 to SANIC 306.

Switch 312 acknowledges the receipt of flits 358 and 364 with acknowledgment flit 370, which is transmitted from switch 312 to SANIC 308. Switch 310 acknowledges the receipt of flits 360 and 366 with acknowledgment flit 372, which is transmitted to switch 312. SANIC 306 acknowledges the
25 receipt of flits 362 and 368 with acknowledgment flit 374 which is transmitted to switch 310.

Architecture Layers and Implementation Overview

A host processor endnode and an I/O adapter endnode typically have
30 quite different capabilities. For example, an example host processor endnode might support four ports, hundreds to thousands of QPs, and allow incoming

RDMA operations, while an attached I/O adapter endnode might only support one or two ports, tens of QPs, and not allow incoming RDMA operations. A low-end attached I/O adapter alternatively can employ software to handle much of the network and transport layer functionality which is performed in hardware (e.g., by SANIC hardware) at the host processor endnode.

One embodiment of a layered architecture for implementing the present invention is generally illustrated at 400 in diagram form in Figure 11. The layered architecture diagram of Figure 11 shows the various layers of data communication paths, and organization of data and control information passed between layers.

Host SANIC endnode layers are generally indicated at 402. The host SANIC endnode layers 402 include an upper layer protocol 404; a transport layer 406; a network layer 408; a link layer 410; and a physical layer 412.

Switch or router layers are generally indicated at 414. Switch or router layers 414 include a network layer 416; a link layer 418; and a physical layer 420.

I/O adapter endnode layers are generally indicated at 422. I/O adapter endnode layers 422 include an upper layer protocol 424; a transport layer 426; a network layer 428; a link layer 430; and a physical layer 432.

The layered architecture 400 generally follows an outline of a classical communication stack. The upper layer protocols employ verbs to create messages at the transport layers. The transport layers pass messages to the network layers. The network layers pass frames down to the link layers. The link layers pass flits through physical layers. The physical layers send bits or groups of bits to other physical layers. Similarly, the link layers pass flits to other link layers, and don't have visibility to how the physical layer bit transmission is actually accomplished. The network layers only handle frame routing, without visibility to segmentation and reassembly of frames into flits or transmission between link layers.

Bits or groups of bits are passed between physical layers via links 434. Links 434 can be implemented with printed circuit copper traces, copper cable, optical cable, or with other suitable links.

The upper layer protocol layers are applications or processes which
5 employ the other layers for communicating between endnodes.

The transport layers provide end-to-end message movement. In one embodiment, the transport layers provide four types of transport services as described above which are reliable connection service; reliable datagram service; unreliable datagram service; and raw datagram service.

10 The network layers perform frame routing through a subnet or multiple subnets to destination endnodes.

The link layers perform flow-controlled, error controlled, and prioritized frame delivery across links.

The physical layers perform technology-dependent bit transmission and
15 reassembly into flits.

Memory Management

A computer system is illustrated generally at 500 in Figure 12.

Computer system 500 includes at least one processor, such as processor 502, for
20 performing sequences of logical operations. Computer system 500 also includes memory 504 for storing instructions and data for use by processor 502. An operating system 506 is stored in memory 504 and controls processor 502 and memory 504 for system operations and for executing processes, such as user process 508 stored in memory 504. Memory 504 typically includes random
25 access memory (RAM), non-volatile memory, and a hard disk drive, but can include any known type of memory storage. Memory 504 can be located local to processor 502 or remote from processor 502. In addition, memory 504 can be shared among numerous other processors.

A consumer process, which can be a kernel process or a user process,
30 such as user process 508, operates under an address space which is pointed to by an effective address of the consumer process. For example, user process 508

operates under effective address space 510. For most advanced processors, the effective address is typically a pointer into a data structure that contains a virtual address. The data structure containing the virtual address is typically referred to as a segment table, such as segment table 512 containing a segment table entry
5 514 holding the virtual address. The effective address pointer is typically controlled by an operating system kernel process, such as virtual memory manager (VMM) 518 of kernel process 516. A user process cannot access a virtual address unless the operating system has allowed the user process to access the virtual address. The operating system allows the access of a virtual
10 address by creating a pointer to the virtual address, including the pointer in the segment table, and allowing the user process to employ the pointer as part of an effective address. The segment table is typically unique to a specific process.

An example effective address 530 is generally illustrated in Figure 13. Example effective address 530 contains a pointer to a segment table entry 532
15 (i.e., effective segment ID); a memory page number 534; and a page offset 536 into the memory page. The pointer to the segment table entry is used to extract the segment table entry from the segment table.

An example segment table entry 540 is illustrated generally in Figure 14. Example segment table entry 540 includes an effective segment ID 532'
20 corresponding to pointer 532 of effective address 530; the upper bits of the virtual address 542 that the effective segment ID is mapped to; and other fields 544, such as segment valid bit field 546.

A physical address is the address actually employed to address the physical memory. A virtual address provides a much wider address space than
25 the actual physical address space. In one embodiment, a user process cannot directly access the virtual address space.

An example virtual address 550 is generally illustrated in Figure 15. Example virtual address 550 includes the upper bits of the virtual address 542' obtained from the segment table entry 540 that maps the effective address to the
30 virtual address; memory page number 534' providing a pointer to the memory page; and page offset 536' providing an offset into the memory page. A physical

address is generated by employing the upper virtual address 542' and page number 534' to extract a page table entry from a page table, such as page table entry 522 of page table 520 stored in memory 504 of Figure 1.

An example page table entry 552 is illustrated generally in Figure 16.

- 5 Example page table entry 552 includes the upper virtual address 542"; physical memory page number 534"; and several other fields 554, such as page valid field 556, and page protection field 558.

In one embodiment of a distributed computer system according to the present invention, such as distributed computer system 30 of Figure 1, one
10 memory management technique is employed for local memory access protection and a separate, but related, memory management technique is employed for remote memory access protection.

When computer system 500 is part of a distributed computer system, such as distributed computer system 30 of Figure 1, computer system 500
15 includes a SANIC 524 for sourcing and sinking frames transmitted on SAN fabric 528. A SANIC driver process (SANIC driver) 526 is stored in memory 504 and controls the SANIC hardware in SANIC 524 and performs other SAN fabric interface functions.

The local memory access protection is tied directly to the operating
20 system kernel and hardware platform's VMM process, such as VMM 518. Local memory access protection assures that a consumer process only accesses the effective address regions authorized by the local operating system kernel process, such as kernel process 516. For example, in a distributed computer system, a consumer process residing on one of the host processor nodes cannot
25 use a SANIC driver, such as SANIC driver 526, to gain unauthorized access to any virtual or physical address region.

The remote memory access protection is under the control of a consumer, which can be a kernel process or a user process. The remote memory access protection assures a process executing on a remote node can only access the
30 effective memory regions authorized by the local consumer process which is

authorized through the local memory access protection by the local operating system VMM.

Local Memory Protection

5 Referring to Figure 12, local memory access protection is performed by executing SANIC driver 526 in a privileged mode. SANIC driver 526 is also tied to the processor platform of computer system 500. In addition, VMM 518 of operating system kernel process 516 preferably controls local memory access protection instead of having local memory access protection being directly
10 controlled by SANIC driver 526. Local memory access protection is preferably controlled by VMM 518 for several reasons including the following reasons.

The VMM controlling the local memory access protection allows the segment table 512 contents to be obtained for the calling process. In order to extract the virtual address from the effective address used by the calling process,
15 SANIC driver 526 needs at least a copy to the segment page table. In one embodiment, SANIC driver 526 requires a pointer to the actual segment table 512 employed by VMM 518.

The VMM 518 controlling local memory access protection allows the segment table entries 514 to be maintained current if the segment table entries
20 change. For example, segment table entries 514 can be changed when a process extends the process's memory region. In order to maintain the segment table entries 514 current, the VMM 518 needs to keep SANIC driver 526 informed of all segment table 512 changes, such as when new segment table entries 514 are added to segment table 512. All segment table entries 514 need to contain
25 similar attributes as the operating system VMM 518's segment table, such as an entry valid attribute. In addition, SANIC driver 526 needs a method of obtaining the latest segment table entries 514 in case SANIC 524 caches some of the latest segment table entries.

The VMM controlling local memory access protection also allows page
30 table entries 522 of page table 520 to be maintained current. In order to maintain page table entries 522 current, VMM 518 keeps SANIC driver 526 informed of

all page table 520 changes, such as page swaps in and page swaps out. All page table entries 522 contain similar attributes as the operating system VMM 518's page table, such as entry valid and protection level attributes. In one embodiment, VMM 518 can grant SANIC 524 access to page table 520.

5 SANIC driver 526 can employ similar mechanisms as employed by processor 502 to restricted access by a calling process to only the effective address the process is authorization to access. One embodiment of a suitable mechanism for managing local memory access protection is illustrated generally at 600 in Figure 17. At step 602, an OpenQP verb is executed to open a QP,
10 such as QP 529, to permit a calling process to communicate with a similar process on a remote node. QP 529 includes a send work queue and a receive work queue. When QP 529 is created, SANIC driver 526 and SANIC 524 save a set of state information about the specific QP instance, such as QP state 531. Example information stored in QP state 531 includes the process using the QP;
15 the size of the QP; the type of transfers the calling process desires for the QP; and the like. The information in QP state 531 is created by SANIC driver 526 and is not accessible by the calling process.

At step 604, when QP 529 is opened, SANIC driver 526 issues a call to the local operating system VMM 518 and obtains the segment table 512 for the
20 calling process that issued the OpenQP verb. This call from SANIC driver 526 to VMM 518 obtains the contents of segment table 512 for the calling process and also registers SANIC driver 526 with the local operating system VMM 518. By registering SANIC driver 526 with VMM 518, VMM 518 is able to notify
25 SANIC driver 526 whenever segment table 512 is updated for the calling process. The segment table 512 is then stored as part of QP state 531. For a process that has many QPs open, the common state information can be stored only once for all the QPs the given process has open. In addition, the QP state can be a pointer to the actual operating system segment table.

At step 606, whenever the process passes an effective address to SANIC
30 driver 526, the segment table 512 stored during QP 529 creation time is employed to translate the effective address into a virtual address.

At step 608, the local operating system VMM 518 informs SANIC driver 526 of any update in segment table 512 by executing a SegmentTableUpdate verb. The SegmentTableUpdate verb is executed by VMM 518 whenever an entry in a segment table 512 is added, deleted, or changed.

5 At step 610, SANIC driver 526 updates QP state 531 to reflect the change in the QP context(s) associated with the updated segment table.

In one embodiment, SANIC 524 employs the same segment table as VMM 518. In this embodiment, updates to the segment table employed by VMM 518 are by definition known by SANIC 524 without execution of the
10 SegmentTableUpdate verb.

At step 612, SANIC 524 maps the virtual address into a physical address by executing a MapLocalPage verb. Even though steps 602 – 608 permit SANIC 524 to map an effective address into a physical address, SANIC 524 must also be able to map a virtual address into a physical address, because many
15 operating system kernel programs use virtual addresses instead of effective addresses for I/O operations.

For example, a disk device driver kernel process typically obtains a virtual address instead of an effective address path to the disk device driver by a process (e.g., file system) that called the disk device driver. In one embodiment,
20 the operating system kernel processes, such as a disk device driver, can use a kernel call to translate the virtual address to a physical address and then pass the physical address to SANIC 524. Alternatively, the operating system kernel processes, such as the disk device driver, can pass the virtual address to SANIC 524 and allow SANIC 524 to translate the virtual address to the physical address
25 by employing the same page table as used by the operating system VMM 518.

At step 614, the consumer can then pass an opaque address to the remote node or alternatively pass the actual physical address to the remote node. In either case, the address is protected by a remote key (RKey) as discussed in detail below. The SANIC would make the remote address opaque versus the
30 actual physical address so as to not expose the physical address to a process in the remote node that will use the physical address for remote direct memory

access (RDMA) operations. Exposing the physical address to the remote node, can expose the physical address on any node in the distributed computer system including the local node, because the remote node could use IPC to pass back the physical address to the local node.

5 In one embodiment, SANIC 524 maintains a private page table, which is not accessible by the calling process. In another embodiment, SANIC 524 has access to the operating system VMM's page table. In the embodiment where SANIC 524 maintains its own private page table, each SANIC needs to have a page table. Each SANIC maintaining its own page table is not very efficient
10 when multiple SANICs are used per node. In the embodiment where SANIC 524 has access to the operating system page table, SANIC driver 526 needs access to operating system VMM 518 system calls that are used to swap pages in the page table to allow unpinned DMAs from remote nodes. The local node passes a remote node a RKey and memory region pair as described below, but
15 without pinning the memory region. The remote node later attempts a RDMA to the memory region, but the page is swapped out, and at this point SANIC 524 needs to swap in the missing page table entry 522.

The above-described mechanism 600 for managing local memory access protection does not require any form of explicit process identification, because
20 SANIC driver 526 is effectively employing the same memory protection mechanism as the local operating system VMM 518. One problem with the above-mechanism for managing local memory access protection is that the page table can become as big as the local operating system's VMM, which can be alleviated to some degree by implementing an inverted page table.

25

Remote Memory Protection

A distributed computer system having one embodiment of a mechanism for managing remote memory access protection according to the present invention is illustrated generally at 700 in Figure 18. Distributed computer
30 system 700 includes a host processor node 702 and a host processor node 703 which both contain mechanisms similar to the mechanism 600 illustrated in

Figure 17 to manage local memory access protection, and both contain components which operate similar to the components described above for computer system 500 of Figure 12.

A host processor node 702 includes at least one processor such as processor 704, for performing sequences of logical operations, and a memory 706 for storing instructions and data for use by processor 704. An operating system 708 containing a kernel process 710 is stored in memory 706 and controls processor 704 and memory 706 for system operations and for executing processes. Operating system kernel process 710 includes a VMM 712. Host processor node 702 also includes SANIC 714 for sourcing and sinking frames transmitted on a SAN fabric 722. A SANIC driver 716 stored in memory 706 controls hardware in SANIC 714 and performs other SAN fabric interface functions. Host processor node 702 includes a user process 718 operating under an effective address space 720.

Host processor node 703 includes at least one processor such as processor 724, for performing sequences of logical operations, and a memory 726 for storing instructions and data for use by processor 724. An operating system 728 containing a kernel process 730 is stored in memory 726 and controls processor 724 and memory 726 for system operations and for executing processes. Operating system kernel process 730 includes a VMM 732. Host processor node 703 also includes SANIC 734 for sourcing and sinking frames transmitted on a SAN fabric 722. A SANIC driver 716 stored in memory 706 controls hardware and SANIC 714 and performs other SAN fabric interface functions. Host processor node 703 includes a user process 738 operating under an effective address space 740.

In an example remote key (RKey) binding operation, a consumer user process, such as user process 718 stored in memory 706 of host processor node 702, or a consumer kernel process (e.g., device driver) can function as a calling consumer process for accessing memory space in memory 706 of host processor node 702, such as effective address space 720 of memory 706.

The RKey binding operation is part of an RKey memory protection mechanism. The RKey binding operation binds a key to a contiguous memory address range, which can be addressed by an effective or virtual address. The contiguous memory address range can have an arbitrary size and an arbitrary alignment. In one embodiment, the RKey memory protection mechanism according to the present invention provides byte level granularity memory access protection.

A user calling process, such as user process 718, can only bind a RKey to an effective address pointing to an address space in memory accessible by the user calling process, such as effective address space 720. On the other hand, a kernel calling process can bind a RKey to either an effective address or a virtual address. Binding a RKey to a virtual address is useful for a kernel calling process when the kernel calling process has received an I/O request which only contains a virtual address plus length from another kernel process. For example, binding a RKey to a virtual address is useful for most kernel device drivers.

In an example operation, a kernel process, such as a file system, issues an I/O request to the kernel device driver. The I/O request from the file system kernel process contains a virtual address plus length, and the type of operation to be performed, such as read operation, disk address operation, or other type of operation. The kernel device driver needs to bind the virtual address plus length to a RKey before employing the virtual address for data transfers to/from the file system kernel process.

In an example RKey binding operation, user process 718 of host processor node 702 has access to an effective address 744 pointing to effective address space 720. User process 718 employs a BindRKey verb to associate effective address 744 to a RKey 746. Host processor node 702 creates a message 748 containing effective address 744 and RKey 746 and transmits message 748 to remote host processor node 703. Later, user process 738 of host processor node 703 employs effective address 744 and RKey 746 to perform a remote DMA (RDMA) operation to access effective address space 720 in host processor node 702. To perform the RDMA operation, host processor node 703

creates a message 750 containing effective address 744 and RKey 746 which is transmitted over SAN fabric 722 to host processor node 702.

One example embodiment of an RKey protection memory mechanism according to the present invention employs four verbs which are AllocateRKey verb; DeallocateRKey verb; BindRKey verb; and UnbindRKey verb.

A consumer process employs the AllocateRKey verb to obtain one or more RKeys. The consumer process employs the BindRKey verb to bind a RKey to a contiguous memory address range. The contiguous memory address range can have an arbitrary size and arbitrary alignment (i.e., location). The contiguous memory address range can be specified by an effective address or a virtual address. However, a consumer user calling process can only use an effective address. A consumer kernel calling process can use an effective address or a virtual address to specify the contiguous memory address range.

If the consumer process is authorized to access the contiguous memory address range to which the RKey will be bound, the BindRKey verb completes successively. If the consumer process is not authorized to access the contiguous memory address range, an invalid parameter result is returned.

The BindRKey verb does not require operating system kernel process involvement. Thus, a light-weight efficient remote memory access protection mechanism embodiment employs hardware to directly execute the BindRKey verb and to not involve operating system kernel software.

In one embodiment, a consumer process employs the BindRKey verb with a null value for the address and a zero value for the length of the contiguous memory address range to implement the UnbindRKey verb to unbind a RKey from a contiguous memory address range. If the consumer process is authorized to access the contiguous memory address range to which the RKey is bound, the UnbindRKey verb completes successfully. If the consumer process is not authorized to access the contiguous memory address range to which the RKey is bound, an invalid parameter result is returned.

In one embodiment, a consumer process can reuse RKeys, but a single RKey cannot be used to protect more than a single memory region at a given instant.

RKeys can be used by all QPs associated with the same address space bound to the RKey. For example, if a process has ten QPs associated to it, each QP can support RDMA access with the same RKey.

The RKey can be used by any number of remote nodes. For example, if a process has a single connecting QP, then only the connected process on the one remote node can use the RKey. However, if a process has many connected QPs, the RKey can be used by all the processes having access to the connected QP. A process with a datagram node QP (e.g., a QP supporting either reliable datagram or unreliable datagram) that allows RDMA operations, can accept RDMA operations from any node in the distributed computer system that obtains the RKey.

In one embodiment, a consumer process employs the DeallocateRKey verb to retire a remote key that was previously obtained through an AllocateRKey verb.

If a RKey and corresponding contiguous memory address range bound to the RKey supplied by a remote node do not match the bound RKey and contiguous memory address range in the local node, the remote node is not granted access to the contiguous memory address range. If the contiguous memory address range is invalid (e.g., no segment table entry for the contiguous memory address range), the remote node is not granted access to the contiguous memory address range. If a virtual address was not previously mapped by the operating system through a map local page verb, the SANIC's default error handler process receives an asynchronous error result to handle the resulting page fault condition.

In one embodiment, an RKey is a one-shot validating mechanism such that as soon as the RKey is used for a RDMA operation, the translation for the RKey is disabled.

Although specific embodiments have been illustrated and described herein for purposes of description of the preferred embodiment, it will be appreciated by those of ordinary skill in the art that a wide variety of alternate and/or equivalent implementations calculated to achieve the same purposes may

5 be substituted for the specific embodiments shown and described without departing from the scope of the present invention. Those with skill in the chemical, mechanical, electro-mechanical, electrical, and computer arts will readily appreciate that the present invention may be implemented in a very wide variety of embodiments. This application is intended to cover any adaptations or

10 variations of the preferred embodiments discussed herein. Therefore, it is manifestly intended that this invention be limited only by the claims and the equivalents thereof.


WHAT IS CLAIMED IS:

1. A method of managing memory in a distributed computer system, the method comprising:
 - 5 binding a remote key to a first address representing a contiguous memory address range at a first endnode;
sending the bound remote key and first address from the first endnode to a second endnode on a communication fabric; and
performing a remote direct memory access operation from the second
 - 10 endnode to access the contiguous memory address range including sending the bound remote key and the first address from the second endnode to the first endnode on the communication fabric.



MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM


Abstract of the Disclosure


- 5 A method of managing memory in a distributed computer system includes binding a remote key to a first address representing a contiguous memory address range at a first endnode. The bound remote key and first address are sent from the first endnode to a second endnode on a communication fabric. A remote direct memory access operation is performed from the second
- 10 endnode to access the contiguous memory address range including sending the bound remote key and the first address from the second endnode to the first endnode on the communication fabric.




2/

**Fig. 2****Fig. 3**



*Fig. 4*

*Fig. 5*

5/

**Fig. 6**

6/

*Fig. 7**Fig. 8*




Fig. 9A






Fig. 9B

**Fig. 10A****Fig. 10B**

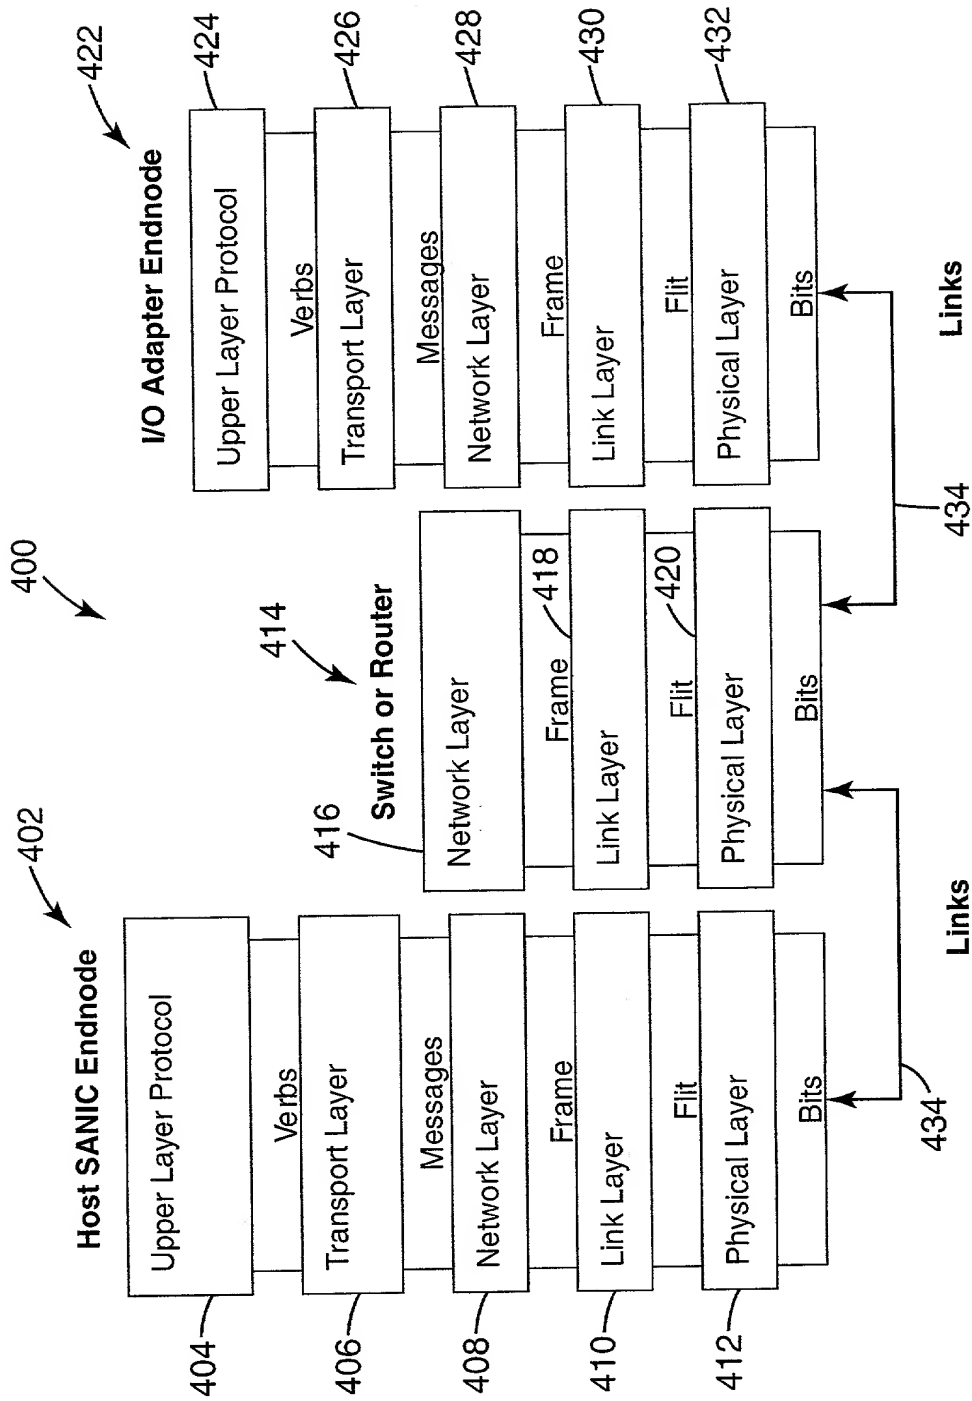


Fig. 11




FIG. 12




FIG. 13




FIG. 14

550



VIRTUAL ADDRESS




FIG. 15

552



PAGE TABLE ENTRY




FIG. 16

09/07/20 09:09:50

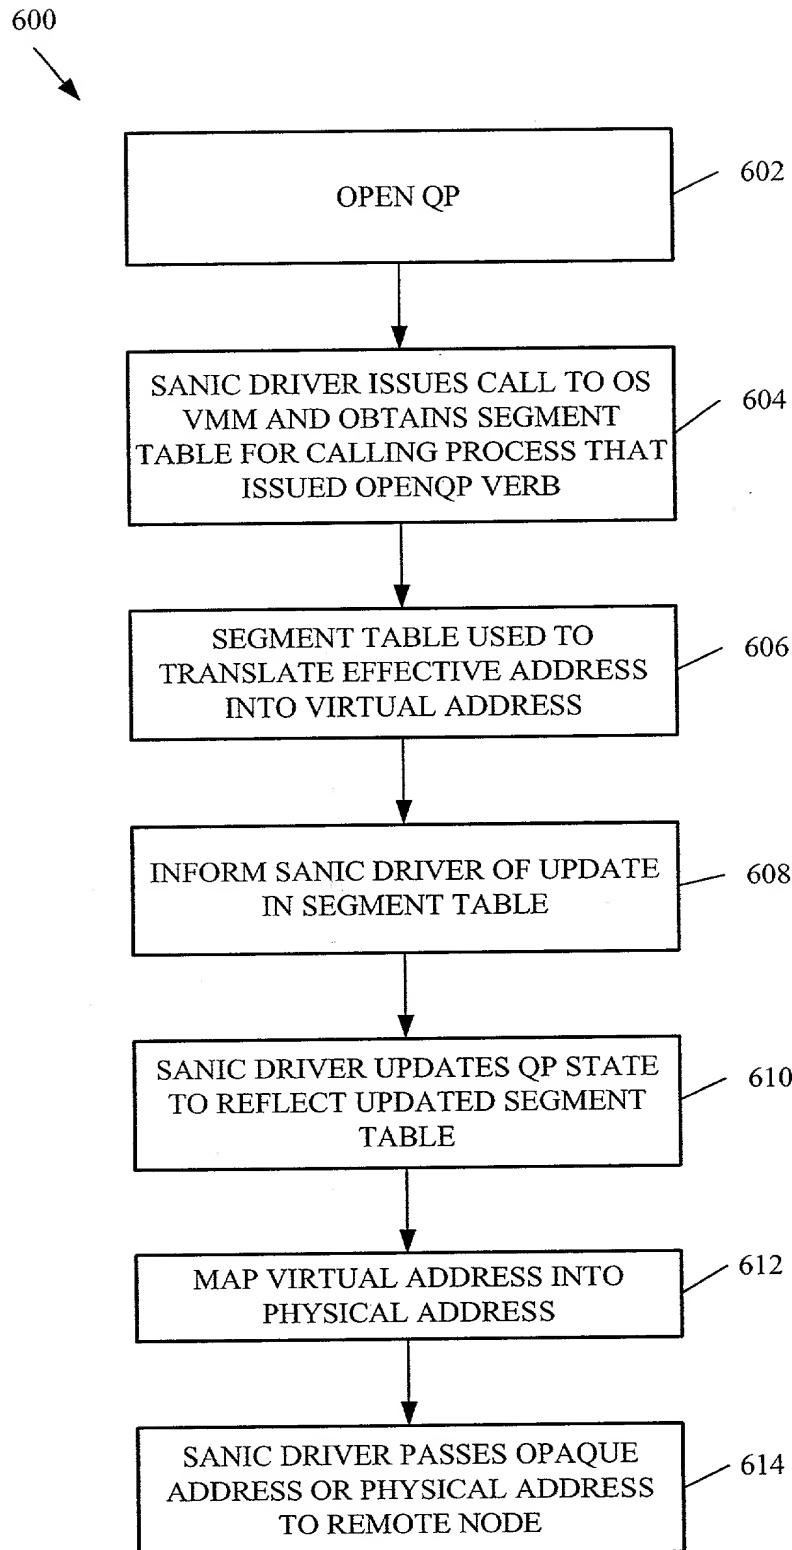


FIG. 17

FIG. 18




FIG. 18

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATIONATTORNEY DOCKET NO. 10002166

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

the specification of which is attached hereto unless the following box is checked:

(X) was filed on Nov. 26, 2001 as US Application No. or PCT International Application
Number 09/980,920 and was amended on 11/26/01 (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119	
PCT	PCT/US00/14282	May 24, 2000	YES: <input checked="" type="checkbox"/>	NO: <input type="checkbox"/>
			YES: <input type="checkbox"/>	NO: <input type="checkbox"/>

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE
60/135,664	5/24/99
60/154,150	9/15/99

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400**Direct Telephone Calls To:****Patrick G. Billig**
(612) 573-2003

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Michael R. KrauseCitizenship: U.S.A.Residence: 220 Hawk Ridge Road, Boulder Creek, CA 95006Post Office Address: Same as above

Inventor's Signature

Date

3/25/2002

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)**

ATTORNEY DOCKET NO. 10002166

Full Name of # 2 joint inventor: Alan F. Benner Citizenship: U.S.A.

Residence: 19 Ridgeview Road, Poughkeepsie, NY 12603

Post Office Address: Same as above

Inventor's Signature _____ Date _____

Full Name of # 3 joint inventor: David J. Garcia Citizenship: U.S.A.

Residence: 24100 Hutchinson Road, Los Gatos, CA 95033

Post Office Address: Same as above

Inventor's Signature _____ Date _____

Full Name of # 4 joint inventor: Renato J. Recio Citizenship: U.S.A.

Residence: 11400 Burnette Road, Austin, TX 78758

Post Office Address: Same as above

Inventor's Signature _____ Date _____

Full Name of # 5 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**ATTORNEY DOCKET NO. **10002166**

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

the specification of which is attached hereto unless the following box is checked:

(X) was filed on **Nov. 26, 2001** as US Application No. or PCT International Application
Number **09/980,920** and was amended on **11/26/01** (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
PCT	PCT/US00/14282	May 24, 2000	YES: X NO: <input type="checkbox"/>
			YES: <input type="checkbox"/> NO: <input type="checkbox"/>

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE
60/135,664	5/24/99
60/154,150	9/15/99

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number **022879**Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400**Direct Telephone Calls To:****Patrick G. Billig**
(612) 573-2003

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: **Michael R. Krause****Citizenship:** **U.S.A.****Residence:** **220 Hawk Ridge Road, Boulder Creek, CA 95006****Post Office Address:** **Same as above**

Inventor's Signature

Date

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)**

ATTORNEY DOCKET NO. 10002166

Full Name of # 2 joint inventor: Alan F. Benner

Citizenship: U.S.A.

Residence: 19 Ridgeview Road, Poughkeepsie, NY 12603 NY

Post Office Address: Same as above

Inventor's Signature Alan F. Benner

Date March 25, 2002

Full Name of # 3 joint inventor: David J. Garcia

Citizenship: U.S.A

Residence: 24100 Hutchinson Road, Los Gatos, CA 95033

Post Office Address: Same as above

Inventor's Signature _____

Date _____

Full Name of # 4 joint inventor: Renato J. Recio

Citizenship: U.S.A.

Residence: 11400 Burnette Road, Austin, TX 78758

Post Office Address: Same as above

Inventor's Signature _____

Date _____

Full Name of # 5 joint inventor: _____

Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____

Date _____

Full Name of # 6 joint inventor: _____

Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____

Date _____

Full Name of # 7 joint inventor: _____

Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____

Date _____

Full Name of # 8 joint inventor: _____

Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____

Date _____

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**ATTORNEY DOCKET NO. 10002166

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

the specification of which is attached hereto unless the following box is checked:

(X) was filed on Nov. 26, 2001 as US Application No. or PCT International Application
Number 09/980,920 and was amended on 11/26/01 (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
PCT	PCT/US00/14282	May 24, 2000	YES: <u>X</u> NO: <u> </u>
			YES: <u> </u> NO: <u> </u>

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE
60/135,664	5/24/99
60/154,150	9/15/99

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400**Direct Telephone Calls To:**Patrick G. Billig
(612) 573-2003

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Michael R. KrauseCitizenship: U.S.A.Residence: 220 Hawk Ridge Road, Boulder Creek, CA 95006Post Office Address: Same as above

Inventor's Signature _____

Date _____

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)

ATTORNEY DOCKET NO. 10002166

Full Name of # 2 joint inventor: Alan F. Benner Citizenship: U.S.A.
Residence: 19 Ridgeview Road, Poughkeepsie, NY 12603
Post Office Address: Same as above
Inventor's Signature _____ Date _____

300
Full Name of # 3 joint inventor: David J. Garcia Citizenship: U.S.A.
Residence: 24100 Hutchinson Road, Los Gatos, CA 95033 CA
Post Office Address: Same as above
Inventor's Signature David J. Garcia Date 3/26/02

Full Name of # 4 joint inventor: Renato J. Recio Citizenship: U.S.A.
Residence: 11400 Burnette Road, Austin, TX 78758
Post Office Address: Same as above
Inventor's Signature _____ Date _____

Full Name of # 5 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____
Residence: _____
Post Office Address: _____
Inventor's Signature _____ Date _____

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATIONATTORNEY DOCKET NO. 10002166

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

MEMORY MANAGEMENT IN DISTRIBUTED COMPUTER SYSTEM

the specification of which is attached hereto unless the following box is checked:

(X) was filed on Nov. 26, 2001 as US Application No. or PCT International Application
Number 09/980,920 and was amended on 11/26/01 (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
PCT	PCT/US00/14282	May 24, 2000	YES: <input checked="" type="checkbox"/> NO: <input type="checkbox"/>
			YES: <input type="checkbox"/> NO: <input type="checkbox"/>

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION NUMBER	FILING DATE
60/135,664	5/24/99
60/154,150	9/15/99

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400**Direct Telephone Calls To:****Patrick G. Billig**
(612) 573-2003

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Michael R. KrauseCitizenship: U.S.A.Residence: 220 Hawk Ridge Road, Boulder Creek, CA 95006Post Office Address: Same as above

Inventor's Signature

Date

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)

ATTORNEY DOCKET NO. 10002166

Full Name of # 2 joint inventor: Alan F. Benner Citizenship: U.S.A.

Residence: 19 Ridgeview Road, Poughkeepsie, NY 12603

Post Office Address: Same as above

Inventor's Signature _____ Date _____

Full Name of # 3 joint inventor: David J. Garcia Citizenship: U.S.A.

Residence: 24100 Hutchinson Road, Los Gatos, CA 95033

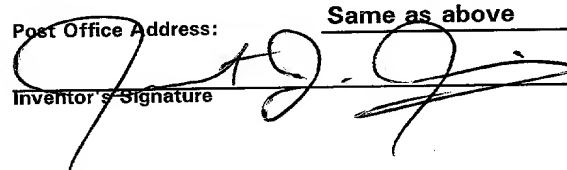
Post Office Address: Same as above

Inventor's Signature _____ Date _____

Full Name of # 4 joint inventor: Renato J. Recio Citizenship: U.S.A.

Residence: 11400 Burnette Road, Austin, TX 78758

Post Office Address: Same as above

Inventor's Signature  Date 3/25/02

Full Name of # 5 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____